

CLAIMS

1. A distributed processing environment, comprising:
 - a client computer platform;
 - a plurality of server computer platforms, each one of said plurality of server computer platforms having an instance of an application residing thereon;
 - a messaging system for controlling the exchange of messages between said client computer platform and said plurality of server computer platforms, said messaging system including messaging services residing at said client computer platform and each one of said plurality of server computer platforms; and
 - a time-distributed load balancing system residing at said client computer platform, said time-distributed load balancing system determining, for a message to be transferred to said application by said messaging system, which one of said plurality of server computer platforms is to receive said message based upon a time associated with said message and at least one message distribution rule requiring transfer of said message to a selected one of said plurality of server computer platforms if said time associated with said message falls within a corresponding one of a plurality of pre-determined time spans.
2. The distributed processing environment of claim 1, wherein said time associated with said message is an arrival time for said message at said messaging service residing at said client computer platform.

3. The distributed processing environment of claim 1, wherein:
said time associated with said message includes a value for a first unit of time and a value
for a second unit of time; and
each one of said plurality of pre-determined time spans encompassing a contiguous range
of values for said second unit of time.

4. The distributed processing environment of claim 3, wherein:
using said at least one message distribution rule, said time-distributed load balancing
system divides said first unit of time into said plurality of pre-determined time spans and
associates each one of said plurality of pre-determined time spans to said corresponding one of
said plurality of server computer platforms; and
said time-distributed load balancing system determining which of said plurality of server
computer platforms is to receive said message based upon which, of said plurality of
pre-determined time spans, said value of said second unit of time associated with said message
falls within.

5. The distributed processing environment of claim 4, wherein said first unit of time is
minutes and said second unit of time is seconds.

6. The distributed processing environment of claim 5, wherein said time associated with
said message is an arrival time for said message at said messaging service residing at said client
computer platform.

7. The distributed processing environment of claim 4, wherein said time-distributed load balancing system residing at said client computer platform further comprises:

a data file for maintaining a list of said plurality of server computer platforms, said pre-determined time span associated with each one of said plurality of server computer platforms and said contiguous range of values for said second unit of time encompassed by said pre-determined time span associated with each one of said plurality of server computer platforms; and

a software application which implements said at least one message distribution rule by determining which one of said plurality of server computer platforms is to receive said message by comparing said value of said second unit of time associated with said message to said contiguous range of values encompassed by each one of said plurality of pre-determined time spans maintained in said data file.

8. The distributed processing environment of claim 7, wherein said software application resides within said messaging service residing at said client computer platform.

9. The distributed processing environment of claim 7 wherein:

said client computer platform is a mainframe computer system;

each one of said plurality of server computer platforms is a mid-range server computer system; and

said messaging services residing at said client computer platform and each one of said server computer platforms are asynchronous messaging services which enable the exchange of messages between a mainframe computer systems and a mid-range server computer system.

10. The distributed processing environment of claim 9, wherein said first unit of time is one minute and said second unit of time is seconds.

11. The distributed processing environment of claim 10, wherein said plurality of predetermined time spans further comprises:

- a first time span which extends between 0 and 20 seconds;
- a second time span which extends between 20 and 40 seconds; and
- a third time span which extends between 40 and 60 seconds.

12. The distributed processing environment of claim 10, wherein said time associated with said message is an arrival time for said message at said messaging service for said client computer platform.

13. The distributed processing environment of claim 12, wherein said software application resides within said messaging service residing at said client computer platform.

14. A distributed processing environment, comprising:

- a client computer platform;
- a plurality of server computer platforms, each one of said plurality of server computer platforms having an instance of an application residing thereon;
- a messaging system for controlling the exchange of messages between said client computer platform and said plurality of server computer platforms, said messaging system including messaging services residing at said client computer platform and each one of said plurality of server computer platforms;
- a time-distributed load balancing and failover system residing at said client computer platform, said time-distributed load balancing and failover system determining, for a message to be transferred to said application by said messaging system, which one of said plurality of server computer platforms is to receive said message based upon a time associated with said message, at least one message distribution rule requiring transfer of said message to a selected one of said plurality of server computer platforms if said time associated with said message falls within a corresponding one of a plurality of pre-determined time spans and at least one failover rule requiring transfer of said message to a subsequent one of said plurality of server computer platforms associated with a subsequent one of said plurality of pre-determined time spans if said selected server computer platform has failed.

15. The distributed processing environment of claim 14, wherein:

 said time associated with said message includes a value for a first unit of time and a value for a second unit of time; and

each one of said plurality of pre-determined time spans encompassing a contiguous range of values for said second unit of time.

16. The distributed processing environment of claim 15, wherein:

using said at least one message distribution rule, said time-distributed load balancing and failover system divides said first unit of time into said plurality of pre-determined time spans, associates each one of said plurality of pre-determined time spans with a corresponding one of said server computer platforms and arranges said plurality of pre-determined time spans into a self-repeating sequence;

said time distributed load balancing and failover system determining which of said plurality of server computer platforms is to receive said message based upon which one, of said plurality of pre-determined time spans, said value of said second unit of time associated with said message falls within;

using said at least one failover rule, if said server computer platform associated with said pre-determined time span within which said value of said second unit of time associated with said message falls within has failed, said time distributed load balancing and failover system selecting, as said subsequent server computer platform to receive said message, said server computer platform associated with a next pre-determined time span in said self-repeating sequence of said pre-determined time spans.

17. The distributed processing environment of claim 16, wherein said first unit of time is one minute and said second unit of time is seconds.

18. The distributed processing environment of claim 17, wherein said time associated with said message is the arrival time for said message at said messaging service residing at said client computer platform.

19. The distributed processing environment of claim 16, wherein said time-distributed load balancing and failover system residing at said client computer platform further comprises:

 a data file for maintaining a list of said plurality of server computer platforms, said predetermined time span associated with each one of said plurality of server computer platforms, said contiguous range of values for said second unit of time encompassed by said pre-determined time span associated with each one of said plurality of server computer platforms and a mark indicating which ones of said plurality of server computer platforms have failed; and

 a software application which implements said at least one message distribution rule by determining which one of said plurality of server computer platforms is to receive said message by comparing said value of said second unit of said time associated with said message to said contiguous range of values encompassed by each one of said plurality of pre-determined time spans maintained in said data file.

20. The distributed processing environment of claim 19, wherein said software application implements said at least one failover rule by determining that, if said server computer platform associated with said pre-determined time span within which said value of said second unit of time associated with said message falls within has failed, said server computer platform associated with said next pre-determined time span in said self-repeating sequence of said pre-determined time spans is to receive said message.

21. The distributed processing environment of claim 20, wherein said first unit of time is one minute and said second unit of time is seconds.

22. The distributed processing environment of claim 21, wherein said time associated with said message is the arrival time for said message at said messaging service residing at said client computer platform.

23. A distributed processing environment, comprising:

- a client computer platform;
- a plurality of server computer platforms coupled to said client computer platform;
- an asynchronous messaging system for controlling the exchange of messages between said client computer platform and said plurality of server computer platforms, said asynchronous messaging system including asynchronous messaging services residing at said client computer platform and each one of said plurality of server computer platforms;
- at least two instances of each one of a plurality of applications distributed among said plurality of server computer platforms, said at least two instances of each one of said plurality of applications residing at respective ones of said plurality of server computer platforms;
- a plurality of data files residing at said client computer platform, each one of said plurality of data files corresponding to one of said plurality of applications;
- each one of said plurality of data files maintaining:
 - a list of each one of said at least two instances of said corresponding one of said plurality of applications and said respective ones of said plurality of server computer platforms on which they reside; and
 - a pre-determined time span assigned to each one of said at least two instances of said corresponding one of said plurality of applications and said respective ones of said server computer platforms on which they reside, for each of said respective ones of said plurality of server computer platforms, said pre-determined time spans assigned to said corresponding one of said at least two instances of said corresponding one of said plurality of applications arranged in a self-repeating sequence, collectively comprise a first unit of time and respectively encompass a contiguous range of values for a second unit of time; and

a load balancing and failover software application residing at said client computer platform;

for each message to be delivered to a first one of said plurality of applications distributed among said plurality of server computer platforms, said load balancing and failover software application:

determining an arrival time of said message at said asynchronous messaging service of said client computer platform, said arrival time of said message including a first value for said first unit of time and a second value for said second unit of time; and

selecting, as said server computer platform to receive said message, said server computer platform corresponding to said pre-determined span of time assigned to said instance of said application for which said second value for said arrival time falls within said contiguous range of values.

24. The distributed computing environment of claim 23, wherein:

each one of said data files further maintains an indication as to whether each one of said at least two instances of said corresponding one of said plurality of applications has failed;

for each message to be delivered to said first one of said distributed applications, said load balancing and failover software application:

determining, from said list, if said selected server computer platform has failed; and

if said selected server computer platform has failed, selecting, as a substitute server computer platform, said server computer platform associated with a next pre-determined time span in said self-repeating sequence of said pre-determined time spans;

repeating said determining and selecting actions until said selected computer platform is not determined to have failed or until all of said server computer platforms have been selected and determined to have failed.

25. The distributed processing environment of claim 24, wherein, for each re-registration of one of said plurality of server computer platforms with said asynchronous messaging system residing at said client computer platform, said load balancing and failover software application removing said failure indication from each of said listed instances of said plurality of applications residing on said re-registering server computer platform.

26. The distributed processing environment of claim 25, and further comprising:

a synchronous messaging service residing on each one of said plurality of server computer platforms, said synchronous messaging service handling the exchange of messages between said asynchronous messaging service residing on said server computer platform and said instances residing on said server computer platform;

said asynchronous messaging service reporting first and second types of responses from said server computer platform selected to receive said message, said first type of message confirming receipt of said message by said messaging service at said selected server and said second type of message confirming receipt of a response to said message;

said load balancing and failover software application:

determining, upon failing to receive said first type of message within a pre-selected time period, that all said instances of said plurality of applications residing at said selected server computer system have failed; and

determining, upon failing to receive said second type of message within said pre-selected time period, that said selected instance of said application residing at said selected server computer system has failed.

27. The distributed processing environment of claim 25, and further comprising:

a synchronous messaging service residing on each one of said plurality of server computer platforms, said synchronous messaging service handling the exchange of messages between said asynchronous messaging service residing on said server computer platform and said instances residing on said server computer platform;

 said asynchronous messaging service reporting first and second types of responses from said server computer platform selected to receive said message, said first type of message confirming receipt of said message by said messaging service at said selected server and said second type of message confirming receipt of a response to said message;

 said load balancing and failover software application:

 determining, upon failing to receive said first type of message within a pre-selected time period, that all said instances of said plurality of applications residing at said selected server computer system have failed; and

 determining, upon failing to receive said second type of message within said pre-selected time period, that all said instances of said plurality of applications residing at said selected server computer system have failed.